

DrIFT: Autonomous Drone Dataset with Integrated Real and Synthetic Data, Flexible Views, and Transformed Domains

Fardad Dadboud¹, Hamid Azad¹, Varun Mehta², Miodrag Bolic¹, Iraj Mantegh²

¹University of Ottawa, ²National Research Council Canada

Abstract

Dependable visual drone detection is crucial for the secure integration of drones into the airspace. However, drone detection accuracy is significantly affected by domain shifts due to environmental changes, varied points of view, and background shifts. To address these challenges, we present the DrIFT dataset, specifically developed for visual drone detection under domain shifts. DrIFT includes fourteen distinct domains, each characterized by shifts in point of view, synthetic-to-real data, season, and adverse weather. DrIFT uniquely emphasizes background shift by providing background segmentation maps to enable background-wise metrics and evaluation. Our new uncertainty estimation metric, MCDO-map, features lower postprocessing complexity, surpassing traditional methods. We use the MCDO-map in our uncertainty-aware unsupervised domain adaptation method, demonstrating superior performance to SOTA unsupervised domain adaptation techniques. The dataset is available at: <https://github.com/CARG-uOttawa/DrIFT.git>.

1. Introduction

Uncrewed Aerial Vehicles (UAVs), also known as drones, have gained popularity in recent years due to their versatility and cost-effectiveness for various operations [18, 62], including healthcare [42], surveillance [10, 16], delivery [36], agriculture [3, 5], construction and mining [9, 54], infrastructure inspection [2], and search-and-rescue [35]. However, their ubiquitous use has raised safety concerns, such as the possibility of their use for malicious activities and collisions with other objects in the airspace [28, 58]. Achieving autonomous flight capabilities in challenging environments, for individual and swarm drones, is vital for various applications [25]. Ensuring the safety of such operations depends on the accurate and efficient processing of drone-related data.

In particular, vision-based drone detection plays a crucial role, as it faces challenges such as distant small objects, handling complex backgrounds (BGs), and distinguishing

drones from other visually similar flying objects. Deep Neural Networks (DNNs) have demonstrated exceptional capabilities in multiple applications, including drone detection [7, 11, 15, 26]. However, distribution shifts from the training to the test set, caused by environmental variations, various points of view (PoVs), and background changes, pose intrinsic challenges in drone detection and affect the DNNs capabilities. Specifically, BG shift, *e.g.* training with data mostly captured with sky background while sky, tree, and ground backgrounds appear in the validation set, is also called unseen BG [63]. Gathering supervised data for all domains to ensure DNN generalization is impractical and often costly for data collection and annotation [33, 37, 46], especially, it is worsened in adverse weather conditions or under regulatory constraints for drone-based applications.

Unsupervised Domain Adaptation (UDA) [6, 8, 21, 37, 39] is a principal approach to addressing domain shift (DS) in object detection (OD). Domain shift refers to shifts in the input image due to environmental factors that affect the performance of drone detection due to their impact on the drone’s appearance in the scene. UDA aims to transfer knowledge from the source to the target domain, despite the lack of supervision in the target domain. This approach has gained popularity in applications such as autonomous vehicles [19, 50, 51, 56] and other edge-AI, where DSs are common, and supervised data is not guaranteed [13, 38]. UDA methods have also been employed extensively to address DSs in drone detection [48, 49]. However, unforeseen situations that cause DSs, such as drones with novel shapes, can still occur. Despite the trend toward using UDA in the field, there is a lack of comprehensive exploration of specialized domain shift and UDA methodologies in drone detection [7, 26, 45, 47, 48, 55, 57, 60, 64, 66]. This gap has catalyzed our work to design a new dataset that addresses these specific challenges.

Combining existing datasets often results in multiple uncontrolled DSs co-occurring, making it difficult to isolate and examine the impact of specific shifts. Moreover, existing datasets lack systematic background segmentation and comprehensive coverage of DS types, making manual annotations costly and infeasible. To overcome this, the DrIFT

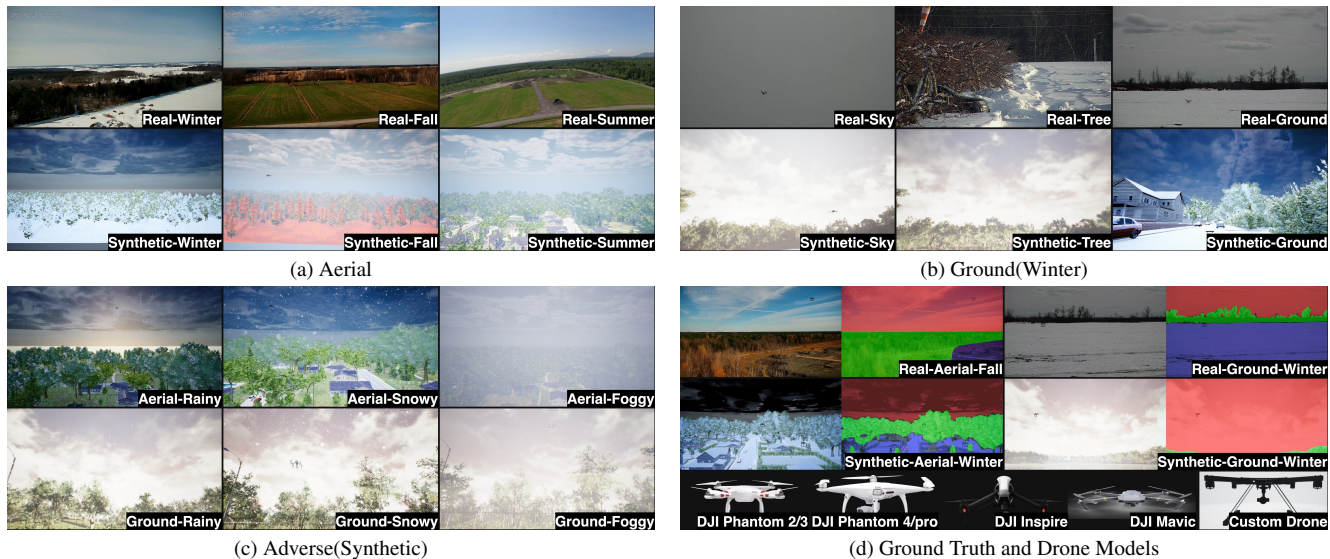


Figure 1. Samples of the DrIFT dataset: a) aerial PoVs in different seasons for both real and synthetic data, b) ground PoV real and synthetic data that has been recorded in winter with a sky, tree, or ground background, c) adverse weather in a synthetic environment for both aerial and ground PoV data, and d) ground truth bounding boxes and background segmentation maps have been illustrated in first two rows. The utilized drones have been also depicted in the last row.

dataset was designed to provide a controlled environment in which individual DS types, PoV, season, weather, and background can be studied independently. DrIFT ensures a balanced distribution across its fourteen distinct domains, addressing limitations in previous datasets and enabling the systematic study of multiple DS types simultaneously. Driven by the need to address these shortages, we present the DrIFT dataset with the following pivotal contributions:

1. The DrIFT dataset introduces a vision-based drone detection dataset. Uniquely, DrIFT comprises fourteen distinct domains constructed by combinations of four major domain shifts: PoV, synthetic-to-real, season, and weather. In most domains, there are sky, trees, and ground backgrounds.
2. We employ BG segmentation maps to introduce the concept of BG shift as a distinct challenge. This novel approach allows us to report BG-wise metrics (*e.g.*, $AP_{50}^S:AP_{50}$ of sky background detections), providing a focused study on how BG shift influence the object detection.
3. We introduce a novel uncertainty evaluation method for OD, surpassing existing methods (Tab. 3). Our method, utilizing a score map, offers significant advantages such as lower complexity of postprocessing and superior capability in capturing DS.
4. Our uncertainty-aware UDA method outperforms state-of-the-art (SOTA) UDA methods for drone detection (Tab. 4).

Dataset	Target	Real/Synt.	PoV	#DS
[11]	Multiple	Real	Gr.	0
[55]	Single	Real	Gr.	0
[26]	Single	Real	Gr.	0
[66]	Single	Real	Aerial	0
[47]	Multiple	Real	Aerial	0
[48]	Multiple	Real+Synt.	-	1
[7]	Multiple	Real	Gr.	0
[60]	Multiple	Real	Aerial	0
[57]	Single	Real	Gr.	0
[45]	Single	Real	Gr.	0
[64]	Multiple	Real	Gr.	0
[1]	Multiple	Real	Gr.	0
[4]	Multiple	Synt.	Gr.+Aerial	0
[15]	Single	Real	Gr.	0
DrIFT	Multiple	Real+Synt.	Gr.+Aerial	4

Table 1. Drone Datasets: The DrIFT dataset, as the first drone detection dataset to study four DSs, includes image frames with multiple drones, real and synthetic data, and ground and aerial PoVs. The "#DS" indicates the number of types of DSs studied.

2. Related Work

2.1. Drone Datasets

Drone datasets have recently become publicly available to address the increasing interest in drone detection [27]. Many datasets have significant limitations, as highlighted in Tab. 1. For instance, the dataset in [11] lacks certain weather conditions and uses a stationary camera.

The datasets in [55] and [26] are limited to a single PoV. The dataset in [66] is restricted to partly cloudy and clear weather. Other datasets [7, 47, 60] offer limited diversity in weather and PoV. UAV-200 [48] uses supervised domain adaptation with a fraction of the target domain during training and only examines the synthetic-to-real DS, while DrIFT studies four types of domain shift in a UDA manner. The number of DS types studied is indicated in Tab. 1. The datasets in [57] and [64] feature multiple drone models but lack comprehensive DSs. The datasets in [1] and [15] focus primarily on ground PoV videos with limited weather conditions. [4] lacks real-world domain.

The **DrIFT dataset** introduces fourteen distinct domains constructed by combinations of four major DS elements: PoV, synthetic-to-real, season, and weather, with sky, trees, and ground backgrounds (Fig. 1). DrIFT uniquely emphasizes BG shift as a separate challenge and employs BG segmentation maps to create BG-wise metrics. This comprehensive approach addresses the lack of datasets that study various DSs in drone detection, making DrIFT the first dataset to comprehensively study all four DSs.

2.2. Land Vehicle Datasets

Land vehicle datasets constitute another topic similar to those of drones within the realm of autonomous vehicles. As inspiration for DrIFT, the SHIFT [56] autonomous driving dataset offers DSs across a spectrum of parameters, such as weather conditions, time of day, and density of vehicle and pedestrian, but does not investigate BG shifts.

2.3. Uncertainty Estimation

Uncertainty estimation is crucial for assessing the safety level of autonomous vehicles, especially drones, by effectively dealing with DS. Conventional methods categorize uncertainty in deep learning into aleatoric and epistemic uncertainties. Aleatoric uncertainty arises from data noise, while epistemic uncertainty is due to limited data or domain coverage, which is more relevant to DS [17].

Historically, uncertainty estimation involves sampling-based techniques like Monte Carlo dropout (MCDO) [20], which, although effective in capturing epistemic uncertainty, are computationally intensive due to their iterative nature and postprocessing complexity [22].

To address computational constraints and accurately capture the epistemic uncertainty arising from data gaps, recent studies [43, 44] explore using gradient self-information directly to assess uncertainty. Nevertheless, they do not inherently encompass the true essence of uncertainty.

To address these computational constraints and the lack of a comprehensive sense of uncertainty, we leverage an efficient approach that combines the strengths of MCDO with a simplified postprocessing mechanism. Our method utilizes MCDO to generate uncertainty maps for each detec-

tion, performing multiple inference passes and aggregating these uncertainties into an overall score map [40], which reduces the postprocessing complexity (Sec. 4.1).

2.4. Detection Calibration Error Estimation (D-ECE)

D-ECE is critical for providing accurate confidence assessments in neural networks, especially for safety-critical applications. The calibration error measures the alignment between the predicted confidence and the actual results, helping to assess the reliability of a model [23, 32]. D-ECE extends from classification-based calibration error estimation but applies specifically to detection tasks, focusing on the regression outputs of object detectors. The concept, introduced by [32], addresses unique detection confidence calibration errors. Further details on its calculation are provided in Sec. 4.1.

2.5. Unsupervised Domain Adaptation (UDA)

UDA addresses domain shift by transferring knowledge from a labeled source domain to an unlabeled target domain. UDA for object detection was first introduced by [8].

Many approaches in UDA have been introduced that come with notable limitations. Pseudo-labeling and self-training methods, such as [30] and [29], generate target pseudo-labels, but incorrect labels can propagate errors, especially in complex backgrounds like our application. Image-to-image translation techniques [24, 52] reduce the domain gap by converting source images into the target style, but these often introduce artifacts and require extensive training data to perform well, which is not feasible in our application.

Among the more recent advancements, uncertainty-aware methods have gained attention for their ability to improve domain adaptation by estimating and incorporating prediction uncertainties. These methods, such as [6, 21, 37, 39], leverage uncertainty metrics to focus on areas where domain shifts are most pronounced. Adversarial training, introduced by [8], complements this by aligning feature distributions between domains. Together, these approaches provide a robust mechanism for handling domain shift, focusing on confident regions and learning domain-invariant features to reduce errors and enhance model robustness. The details of our approach are discussed further in Sec. 4.1.

3. DrIFT Dataset

We have developed a vision-based drone detection dataset consisting of image frames, ground truth bounding boxes, and BG segmentation maps (Fig. 1d). In Sec. 3.1, an overview of the DrIFT dataset’s sensor, experimental setup, annotation, and dataset design has been presented. In the following, precise information regarding DrIFT’s various

domains has been compiled to represent the dataset’s purpose for the DS. For more detailed statistics of DrIFT, the reader can go through the supplementary materials.

3.1. The DrIFT Story

Real Ground PoV’s video recordings for the DrIFT dataset were captured with a Bosch pan-tilt-zoom (PTZ) camera. The DJI Phantom 2/3, Phantom 4/Pro, Inspire, and Mavic (Fig. 1d) were captured between 0.1 and 1.5 kilometers away in the recordings. A drone is predominantly present in the frames. The semi-automatic annotation has been done using the CVAT [14]. We have generated multiple other domains of data in our dataset to represent the DS.

Real Aerial PoV has been added to the DrIFT dataset to achieve the PoV shift concept. For the aerial PoV, a custom-built drone model was utilized (Fig. 1d). In this experiment, mobile electro-optical cameras, the Infiniti STR-8MP-3X and GoPro were used to record multiple drone footage between 20 and 100 meters in the line of sight. The frames were recorded in different seasons, resulting in various BGs, such as the sky, trees in various seasons, and the ground with different colors.

Synthetic Data is recorded in the AirSim [53] simulator for simulating real-world data counterparts in a simulated environment for all domains for considering synthetic-to-real domain shift, and due to the impossibility of flying in adverse weather conditions.

It is a common practice for domain-adaptive network training to have the same number of samples in the source and target domains [12]. Therefore, we designed the DrIFT dataset to maintain a balanced number of samples across domains within both the training and validation sets as long as we had sufficient real data for the domains.

Background segmentation, as one of the contributions of the DrIFT dataset, is important for its innovative exploration of BG shift. All validation frames’ backgrounds have been segmented into sky, tree, and ground segments (Fig. 1d) using the Track Anything platform [31, 61]. By utilizing segmentation maps, it becomes feasible to utilize different metrics corresponding to different backgrounds (Tab. 2, details in Sec. 4.1).

All annotations were then double-checked and refined by human annotators to ensure accuracy.

3.2. Dataset Design

To address a deficiency in drone detection datasets, we designed DrIFT with a concentration on studying common domain shifts in the wild.

Synthetic-to-Real: In practical scenarios, capturing every conceivable real-world situation can be infeasible due to logistical challenges, resource limitations, and the prohibitive costs of annotation. To this end, we brought up syn-

thetic data in order to initiate research on synthetic-to-real DS. In DrIFT, all real-world data domains have simulated counterparts except for adverse weather conditions that do not exist in our real-world part of the dataset.

PoV Shift: The camera’s PoV change (ground and aerial) contains different BGs and orientations of the target objects. This shift can significantly impact detection performance, making it a distinct type of DS.

Weather Shift: drones cannot be easily deployed in adverse weather. On the other hand, because this is a common DS in the wild, the system must be robust. Therefore, synthetic data is collected to study weather DS.

Background Shift: The unseen background problem [63], also called BG shift in DrIFT, is present in various drone detection or autonomous driving datasets regarding the aforementioned DS. Nevertheless, no study has explicitly looked into the BG shift in object detection using BG segmentation maps. DrIFT investigates the BG shift from the sky to the tree and ground.

4. DrIFT Benchmark

This section first provides a comprehensive overview of the methodology used for the benchmark. The following subsection provides a comprehensive overview of the different benchmark scenarios. Subsequently, the results of the benchmark are reported. This section concludes with a comprehensive analysis of the benchmark outcomes and the dominant challenges of the DrIFT dataset. The supporting statements will be presented in the supplementary materials.

4.1. Methodology

4.1.1 Problem definition

The primary goal of the DrIFT benchmark is to evaluate the performance of OD models under various shifts and the capabilities of UDA methods to address this issue. Performance metrics include average precision (AP), uncertainty metrics, and D-ECE which are reported BG-wise.

Let $\mathcal{D} = \{(\mathbf{X}_i, \mathbb{Y}_i)\}_{i=1}^N$ be the dataset, where \mathbf{X}_i are the input images and \mathbb{Y}_i are the set of ground truth annotations containing bounding box coordinates and class labels for objects within each input image. The OD model predicts a set of detections $\hat{\mathbb{Y}}_i = \{\hat{\mathbf{y}}_i^j\}_{j=1}^{N_{x_i}}$, where each detection $\hat{\mathbf{y}}_i^j = (\hat{\mathbf{b}}_i^j, \hat{c}_i^j, \hat{s}_i^j)$ consists of bounding box coordinates $\hat{\mathbf{b}}_i^j$, class label \hat{c}_i^j , and confidence score \hat{s}_i^j . N and N_{x_i} are the number of samples in the dataset and the number of detections for the i -th input image, respectively.

In OD, after initial detections, the Non-Maximum Suppression (NMS) process [41] filters out redundant or sub-optimal detections. NMS first generates a set of candidates for each detection. These candidates are defined as all other predictions sharing the same class label and having

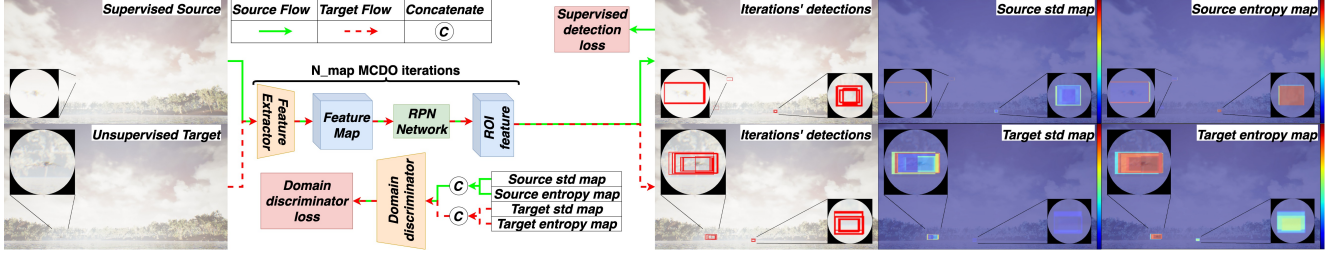


Figure 2. Uncertainty-aware UDA framework: In addition to supervised learning on the source domain, the concatenated std and entropy maps are input into the discriminator as part of the adversarial learning process. Magnified regions around detections are shown for better visualization. Colorbars are placed on the right side of the std and entropy maps. Green-solid and red-dash lines represent the source and target domain paths, respectively. All detections from multiple iterations are displayed to illustrate the generation and behavior of the uncertainty maps.

an Intersection-over-Union (IoU), a measure of the overlap between bounding boxes, above a threshold ϵ . Detections below a confidence threshold δ are then discarded. The candidate set for a given detection \hat{y}_i^j is defined as:

$$\mathbb{C}_{\hat{y}_i^j} = \{\hat{y}_i^k | k \neq j, IoU(\hat{y}_i^k, \hat{y}_i^j) \geq \epsilon, \hat{c}_i^k = \hat{c}_i^j, \hat{s}_i^k \geq \delta\}. \quad (1)$$

After filtering, the remaining detection with the highest confidence score is retained as the final prediction. AP [41] has been employed to quantify OD performance.

4.1.2 Performance Metrics

D-ECE [32] is just used in our benchmark to study domain shift impacts on the calibration error. D-ECE [32] was calculated by binning the confidence space as well as box coordination parameters space in which there are N_k equally distributed bins corresponding to the k -th dimension, $N_{total} = \prod_{k=1}^K N_k$. The goal of binning is to account for variations in calibration error across different confidence levels and spatial dimensions, ensuring that errors are captured in an unbiased manner. Therefore, D-ECE could be formalized

$$D-ECE_k = \sum_{n=1}^{N_k} \frac{|I(n)|}{|\hat{y}|} |prec(n) - conf(n)|. \quad (2)$$

Within Eq. (2), $|I(n)|$ is used to describe the cardinality of the bin, whereas $|\hat{y}|$ represents the total number of detections. $conf(n)$ denotes the mean confidence score of the detections within the bin, whereas $prec(n)$ is a statistical metric that quantifies the proportion of true positives among the detections in the bin.

We utilize MCDO-based and gradient self-information metrics to estimate uncertainty in the presence of DS and compare their capabilities with our proposed method, **MCDO-map**, to take advantage of them in our UDA

method. The utilized methods are referred to as **MCDO-NMS** and **Grad-loss**, respectively. The **Grad-loss** captures the degree of epistemic uncertainty for each detection. **Grad-loss-localization** and **Grad-loss-classification** refer to the localization and classification terms, respectively.

The MCDO-based method involves running multiple inference passes with dropout activated. Detections are matched to a candidate list based on the highest IoU threshold. The standard deviation of the localization parameters and the entropy of the mean classification probabilities are calculated for each list. This technique includes **MCDO-NMS-localization** and **MCDO-NMS-classification**. For details on these methods, please refer to the supplementary materials.

As opposed to utilizing NMS-based or data association techniques in an MCDO scheme, a score map is constructed in a pixel-wise manner. Given the predictions, we convert the detection outputs to a 3D map. Let $\hat{Y}_i = \{\hat{y}_i^j\}_{j=1}^{N_{x_i}}$ be the set of detections for an input image. The score map \mathcal{S} is a tensor of shape $(H, W, C + 1)$, where H and W are the height and width of the input image, and C is the number of classes. For each detection \hat{y}_i^j , the score \hat{s}_i^j is assigned to each pixel inside the bounding box \hat{b}_i^j ,

$$\mathcal{S}(x, y, c) += \begin{cases} \hat{s}_i^j & (x, y) \in \hat{b}_i^j, c = \hat{c}_i^j \\ (1 - \hat{s}_i^j) & (x, y) \in \hat{b}_i^j, c = \text{BG} \end{cases} \quad (3)$$

The \mathcal{S} is zero initiated, resulting in all-zero vectors for pixels that are not contained within any bounding box. For these pixels, we replace the all-zero vectors with a vector with a 1 for the background element and zeros for all other elements. After populating the score map, we normalize it using the softmax function. Next, we calculate the mean and standard deviation of the score map over multiple iterations of our object detector forward path that are

$$\bar{\mathcal{S}} = \frac{1}{N_{\text{map}}} \sum_{n=1}^{N_{\text{map}}} \mathcal{S}_n, \sigma_{\mathcal{S}} = \sqrt{\frac{1}{N_{\text{map}}} \sum_{n=1}^{N_{\text{map}}} (\mathcal{S}_n - \bar{\mathcal{S}})^2}. \quad (4)$$

Finally, we compute the entropy of the mean score map, $H_{\mathcal{S}}(x, y, :) = -\sum_c \mathcal{S}(x, y, c) \log \mathcal{S}(x, y, c)$, and concatenate the standard deviations to create the uncertainty map, $\text{MCDO-map} = \text{concat}(H_{\mathcal{S}}(x, y, :), \sum_c \sigma_{\mathcal{S}}(x, y, c))$. From an intuitive standpoint, it can be observed that increasing changes in localization parameters of the predictions are associated with a corresponding increase in the standard deviation of the boundaries surrounding pixels. For example, in Fig. 2, the left magnified detection in the target std map shows higher deviation (with colors closer to red) compared to the source std map, where the corresponding detection is mostly blue, indicating lower deviation across pixels. Similarly, a higher frequency of change in prediction scores is shown to be linked to an elevated level of entropy. The same behavior in the entropy maps can be observed in Fig. 2. In contrast to traditional MCDO-based approaches, instead of handling individual bounding boxes from each iteration and suffering postprocessing complexity [17], our method generates a pixel-wise score map during each iteration and avoids complex postprocessing.

4.1.3 Domain Shift and Adaptation

DS occurs when the training (source) domain \mathcal{D}^s and the testing (target) domain \mathcal{D}^t differ, leading to a performance drop in machine learning models. Let $\mathcal{D}^s = \{(\mathbf{X}_i^s, \mathbf{Y}_i^s)\}_{i=1}^{N^s}$ and $\mathcal{D}^t = \{(\mathbf{X}_i^t, \mathbf{Y}_i^t)\}_{i=1}^{N^t}$. We denote the source and target distributions as $p^s(\mathbf{X}, \mathbf{Y})$ and $p^t(\mathbf{X}, \mathbf{Y})$, respectively. Distribution shift is defined as $p^s(\mathbf{X}, \mathbf{Y}) \neq p^t(\mathbf{X}, \mathbf{Y})$. If we consider $p^s(\mathbf{X}, \mathbf{Y}) = p^s(\mathbf{X})p^s(\mathbf{Y}|\mathbf{X})$ and $p^t(\mathbf{X}, \mathbf{Y}) = p^t(\mathbf{X})p^t(\mathbf{Y}|\mathbf{X})$, the DS happens when $p^s(\mathbf{X}) \neq p^t(\mathbf{X})$, $p^s(\mathbf{Y}|\mathbf{X}) = p^t(\mathbf{Y}|\mathbf{X})$.

Our UDA method focuses on leveraging uncertainty information to enhance the robustness of the object detector in the presence of DS. We got inspired by ADVENT [59] while modifying it by changing the representation of input data to the discriminator and introducing a novel uncertainty estimation method. The intuition behind this method is that DS introduces uncertainty in predictions, especially in regions where the model is less confident. Our uncertainty maps highlight areas where the domain shift has the most impact, guiding the adaptation process to focus on these challenging regions. Following the concatenation process, the **MCDO-map** is subsequently forwarded to a domain discriminator to fool it, initiating adversarial training (Fig. 2). The calculation of the overall loss is $\mathcal{L}_{total} = \mathcal{L}_{detection} - \lambda \times \mathcal{L}_{adv}$, where the detection loss $\mathcal{L}_{detection}$ is a combination of cross-entropy classification and smooth L_1 regression loss, $\mathcal{L}_{detection} = \mathcal{L}_{cls} + \mathcal{L}_{reg}$. The adversarial loss \mathcal{L}_{adv} is

$$\begin{aligned} \mathcal{L}_{adv} = & -\mathbb{E}_{\mathbf{X} \sim p^s} \log D(\text{MCDO-map}(\mathbf{X})) \\ & -\mathbb{E}_{\mathbf{X} \sim p^t} \log(1 - D(\text{MCDO-map}(\mathbf{X}))), \end{aligned} \quad (5)$$

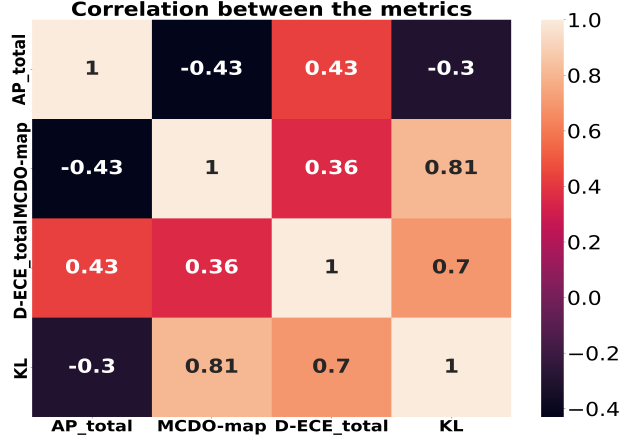


Figure 3. Correlation heatmap: The KL divergence of feature maps distributions and all metrics are calculated between the source (domain I in Tab. 2) and target domains. MCDO-map’s high positive correlation with KL divergence shows high capabilities of MCDO-map to capture the DS.

when the D is the discriminator network. The detection base network is updated to minimize the total loss, $\omega^* = \arg \min_{\omega} \mathcal{L}_{total}$, while the discriminator network is updated to maximize the adversarial loss, $\theta^* = \arg \max_{\theta} \mathcal{L}_{adv}$.

4.2. Benchmark Scenarios

We will begin our benchmark with Tab. 2, illustrating the impact of domain shift on object detection using the AP, uncertainty, and D-ECE metrics. Tab. 3 presents a comparison between our proposed MCDO-map method and other uncertainty metrics. Finally, our novel uncertainty-aware UDA object detector is compared with SOTA UDA methods in Tab. 4. Supplementary materials have been provided to support our discussions.

Background-wise Metrics: To assess OD performance under BG shifts, we introduce BG-wise metrics, which calculate metrics separately for different BGs (e.g., sky, tree, ground). Given detections $\hat{\mathbf{Y}}$ and ground truth \mathbf{Y} , we classify each into BG categories using segmentation maps to identify the background category to which most of the pixels within the bounding box belong. The metric M for each BG category can be expressed as $M_{bg}(\hat{\mathbf{Y}}, \mathbf{Y}) = M(\hat{\mathbf{Y}}_{bg}, \mathbf{Y}_{bg})$, $bg \in \{\text{sky, tree, ground}\}$. These metrics provide a detailed analysis of how different BGs affect object detection performance.

4.3. Experiments and Results

Tab. 2 shows different DS scenarios and their impact on object detection models. It is our contribution that metrics are reported background-wise, highlighting the influence of background shifts. The AP under the Sky column in row

Source Domain	Target Domain				$AP_{50} \uparrow \%$				MCDO-map \downarrow	D-ECE $\downarrow \times 10^{-6}$			
	PoV	Source	Season	Weather	Total	Sky	Tree	Ground	$\times 10^{-4}$	Total	Sky	Tree	Ground
I	ground	synthetic	winter	normal	40.7	67.1	0.2	54.7	3582	674	248	576	1176
	ground	real	winter	normal	3.2	9.5	0.1	0.0	5355	2253	1586	69	2623
	ground	synthetic	adverse	rainy	-	75.8	-	-	3866	-	7679	-	-
	ground	synthetic	adverse	snowy	-	63.0	-	-	4686	-	6803	-	-
	ground	synthetic	adverse	foggy	-	98.7	-	-	4454	-	1796	-	-
	aerial	synthetic	winter	normal	12.7	35.6	0.0	2.4	4287	4095	2063	2819	5518
	aerial	synthetic	fall	normal	41.5	81.0	10.8	32.9	4680	7124	8092	5275	3022
	aerial	synthetic	summer	normal	51.5	74.3	64.8	15.3	4677	6447	7327	2245	1444
	aerial	synthetic	adverse	rainy	-	26.1	-	-	4084	-	2923	-	-
	aerial	synthetic	adverse	snowy	-	19.8	-	-	4835	-	8008	-	-
	aerial	synthetic	adverse	foggy	-	35.7	-	-	4394	-	4142	-	-
II	ground	real	winter	normal	34.4	73.7	0.5	29.0	2673	1242	1105	1694	0329

Table 2. DS impact on Faster R-CNN detector trained on source domain I or II in terms of AP, uncertainty, and D-ECE. The Faster R-CNN detector is validated on different target domains. Shaded cells show the shifted element in each target domain. AP and D-ECE are reported BG-wise, although the MCDO-map uncertainty is reported totally. For adverse weather, we have only the sky background. Source domain I: ground-synthetic-winter-normal-sky, source domain II: ground-real-winter-normal-sky

one is the reference. Significant AP decreases are evident, such as in row two when shifting from synthetic to real. PoV and weather shifts in rows 9 to 11 also show notable changes. Comparing the reference AP and the sky APs in other rows demonstrates decreased APs for tree and ground backgrounds (see Fig. 7 in supplementary material). The domain I in Tab. 2 is the source domain all over the text unless other domains are mentioned.

We calculated the Kullback-Leibler (KL) divergence between source and target domain feature map distributions (see supplementary material) to analyze the relationship between metrics and different shifts from Tab. 2. Fig. 3 shows a heatmap of Pearson correlations among AP, D-ECE, MCDO-map, and KL divergence. The high positive correlation between our MCDO-map and KL divergence indicates the MCDO-map’s effectiveness in capturing DS.

The negative correlation of AP with MCDO-map and KL divergence suggests that higher AP corresponds to lower uncertainty and smaller feature map distribution distances. The positive correlation between AP and D-ECE indicates model miscalibration under DS. Additionally, the positive relation between D-ECE, MCDO-map, and KL divergence highlights their significant association with DS.

The goal of Tab. 3 is to compare our MCDO-map method with the MCDO-NMS and Grad-loss. Our method consistently shows increased uncertainty with DS, highlighting its effectiveness in capturing DSs. The wider violin plots for the MCDO-map in Fig. 4 of supplementary material demonstrate its superior capability to separate different DS levels compared to other metrics. MCDO-NMS-Classification for TPs shows some capability in separat-

ing different DSs, highlighted in Tab. 3 and supplementary material, but requires supervision and often decreases with DS. Grad-loss-localization is consistently capturing the DS (Tab. 3) but lacks the potential to separate DSs effectively.

Our results in Fig. 3 and supplementary material further support the MCDO-map’s effectiveness in capturing DS. Thus, we conclude that the MCDO-map is the best method for our UDA approach, offering significant improvements over traditional uncertainty estimation techniques.

In Tab. 4, the results of some SOTA UDA object detectors on the DrIFT dataset are reported alongside our results. One significant DS is from sky to tree, where AP dropped from 67.1 to 0.2 (Tab. 2, first row). Our UDA method outperforms others with an AP of 10.7 for the tree background and 46.3 in total, demonstrating its effectiveness in adapting to different BGs. Similarly, for the ground background domain, our method achieves an AP of 44.8 in total and 1.2 for the tree background, showcasing robustness.

For the aerial-synthetic-winter-normal domain, multiple BGs in each frame could take place, so we did not specify any BG for it. Our UDA method achieves the highest AP in total (17.8), sky (41.3), and tree (0.5), indicating its capability to adapt to different PoVs and complex scenes. In the ground-real-winter-normal-sky domain, our method achieves the highest AP in total (5.7) and tree (0.8), proving its effectiveness with real-world data and different seasons. The results are consistent for tree and ground backgrounds as well, demonstrating our method’s adaptability for two types of DSs occurring simultaneously.

Our method aims to deceive a domain discriminator by making the uncertainty maps for both source and target do-

Validation Domain				MCDO-NMS $\times 10^{-3}$						MCDO-Map	grad-loss $\times 10^{-3}$	
PoV	Source	Season	Weather	Localization			Classification			$\times 10^{-4}$	Loc.	Cls.
				Total	TP	FP	Total	TP	FP			
ground	synthetic	winter	normal	107	63	107	436	141	439	3582	445	564
ground	real	winter	normal	173	56	180	495	298	508	5355	434	692
ground	synthetic	adverse	rainy	98	67	99	401	112	404	3866	390	634
ground	synthetic	adverse	snowy	79	58	79	454	183	456	4686	458	658
ground	synthetic	adverse	foggy	85	83	85	449	241	484	4454	335	642
aerial	synthetic	winter	normal	101	62	103	407	166	417	4287	444	500
aerial	synthetic	fall	normal	92	105	92	523	312	529	4680	435	653
aerial	synthetic	summer	normal	94	92	94	538	269	543	4677	434	658
aerial	synthetic	adverse	rainy	100	092	101	440	238	446	4084	468	554
aerial	synthetic	adverse	snowy	85	67	85	598	298	605	4835	445	690
aerial	synthetic	adverse	foggy	105	116	103	468	390	477	4394	445	635

Table 3. Uncertainty Metrics Comparison: The network is Faster R-CNN trained on the Tab. 2’s domain I. Each row shows the validation domain in which experiments have been done for the sky background. In each row, the uncertainty level was evaluated using three different methods. MCDO-NMS reported separately for TP and FP detections. Shaded cells show the shifted element. Loc.: localization, Cls.: Classification

Method	Total	Sky	Tree	Gr.	Method	Total	Sky	Tree	Gr.	Method	Total	Sky	Tree	Gr.
S.O.	40.7	67.1	0.2	54.7	S.O.	40.7	67.1	0.2	54.7	S.O.	12.7	35.6	0.0	2.4
C.fMix [34]	44.0	66.8	5.4	56.3	C.fMix [34]	42.5	64.5	0.5	56.0	C.Mix [34]	15.3	37.1	0.4	3.2
SAPN [65]	41.6	61.4	5.1	53.2	SAPN [65]	41.6	61.4	5.1	53.2	SAPN [65]	14.3	35.9	0.3	2.9
PT [6]	42.5	64.9	6.3	55.3	PT [6]	42.0	61.5	0.9	54.2	PT [6]	14.7	36.6	0.6	3.1
Ours	46.3	62.1	10.7	55.9	Ours	44.8	66.0	1.2	56.5	Ours	17.8	41.3	0.5	3.6
(a) ground-synthetic-winter-normal-tree					(b) ground-synthetic-winter-normal-ground					(c) aerial-synthetic-winter-normal				
Method	Total	Sky	Tree	Gr.	Method	Total	Sky	Tree	Gr.	Method	Total	Sky	Tree	Gr.
S.O	3.2	9.5	0.1	0.0	S.O.	3.2	9.5	0.1	0.0	S.O.	3.2	9.5	0.1	0.0
C.Mix [34]	5.0	13.2	0.7	0.3	C.Mix [34]	7.0	10.5	5.0	0.3	C.Mix [34]	7.1	10.0	4.8	0.4
SAPN [65]	4.8	12.8	0.6	0.2	SAPN [65]	6.8	10.0	4.8	0.2	SAPN [65]	6.8	9.8	4.9	0.2
PT [6]	4.6	12.5	0.5	0.2	PT [6]	6.5	10.2	4.6	0.3	PT [6]	7.1	9.7	5.2	0.3
Ours	5.7	14.0	0.8	0.4	Ours	8.2	11.0	6.0	0.5	Ours	8.9	9.9	4.6	0.9
(d) ground-real-winter-normal-sky					(e) ground-real-winter-normal-tree					(f) ground-real-winter-normal-ground				

Table 4. Background-wise mAP of different UDA methods for different domain shifts: Our UDA method surpasses the SOTA techniques in most cases. The target domain is written below each subtable. S.O.: Source Only (trained network on Tab. 2’s domain I).

mains nearly identical. This dual focus on source and target domain alignment is crucial for robust performance across various DSs. The adaptation process involves a trade-off, accepting some degradation in the source domain to achieve significant improvements in the target domain.

5. Conclusion

The DrIFT dataset addresses the need for a vast study of domain shift in drone detection by introducing fourteen distinct domains and emphasizing background shift utilizing

background segmentation maps. Our findings show a positive correlation between MCDO-map uncertainty, domain shift, and D-ECE, and a negative correlation with AP. The MCDO-map outperformed other uncertainty metrics in capturing domain shift in the DrIFT dataset. Our uncertainty-aware UDA on object detection also surpassed SOTA methods in the DrIFT dataset. In the future, we aim to explore more nuanced domain adaptation techniques that minimize the source domain performance degradation, which is a drawback for our UDA method in some cases.

References

- [1] SO Ajakwe, VU Ihekoronye, G Mohtasin, R Akter, A Aouto, DS Kim, and JM Lee. Visiodect dataset: An aerial dataset for scenario-based multi-drone detection and identification. *IEEE Dataport*, 2022. 2, 3
- [2] Abdulla Al-Kaff, Francisco Miguel Moreno, Luis Javier San José, Fernando García, David Martín, Arturo de la Escalera, Alberto Nieva, and José Luis Meana Garcéa. Vbii-uav: Vision-based infrastructure inspection-uav. In *Recent Advances in Information Systems and Technologies: Volume 2* 5, pages 221–231. Springer, 2017. 1
- [3] Paul Albert, Mohamed Saadeldin, Badri Narayanan, Brian Mac Namee, Deirdre Hennessy, Noel E O’Connor, and Kevin McGuinness. Unsupervised domain adaptation and super resolution on drone images for autonomous dry herbage biomass estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1636–1646, 2022. 1
- [4] Antonella Barisic, Frano Petric, and Stjepan Bogdan. Sim2Air - Synthetic Aerial Dataset for UAV Monitoring. *IEEE Robotics and Automation Letters*, 7(2):3757–3764, Apr. 2022. 2, 3
- [5] Widodo Budiharto, Andry Chowanda, Alexander Agung Santoso Gunawan, Edy Irwansyah, and Jarot Sembodo Suroso. A review and progress of research on autonomous drone in agriculture, delivering items and geographical information systems (gis). In *2019 2nd world symposium on communication engineering (WSCE)*, pages 205–209. IEEE, 2019. 1
- [6] Meilin Chen, Weijie Chen, Shicai Yang, Jie Song, Xinchao Wang, Lei Zhang, Yunfeng Yan, Donglian Qi, Yueting Zhuang, Di Xie, et al. Learning domain adaptive object detection with probabilistic teacher. *arXiv preprint arXiv:2206.06293*, 2022. 1, 3, 8
- [7] Yueru Chen, Pranav Aggarwal, Jongmoo Choi, and C-C Jay Kuo. A deep learning approach to drone monitoring. In *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 686–691. IEEE, 2017. 1, 2, 3
- [8] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. 1, 3
- [9] Hee-Wook Choi, Hyung-Jin Kim, Sung-Keun Kim, and Wongi S. Na. An overview of drone applications in the construction industry. *Drones*, 7(8), 2023. 1
- [10] Jinwoo Choi, Gaurav Sharma, Manmohan Chandraker, and Jia-Bin Huang. Unsupervised and semi-supervised domain adaptation for action recognition from drones. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1717–1726, 2020. 1
- [11] Angelo Coluccia, Alessio Fascista, Arne Schumann, Lars Sommer, Anastasios Dimou, Dimitrios Zarpalas, Miguel Méndez, David De la Iglesia, Iago González, Jean-Philippe Mercier, et al. Drone vs. bird detection: Deep learning algorithms and results from a grand challenge. *Sensors*, 21(8):2824, 2021. 1, 2
- [12] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [13] Gabriela Csurka, Timothy M Hospedales, Mathieu Salzmann, and Tatiana Tommasi. Visual domain adaptation in the deep learning era. *Synthesis Lectures on Computer Vision*, 11(1):1–190, 2022. 1
- [14] CVAT.ai Corporation. Computer Vision Annotation Tool (CVAT), 9 2022. 4
- [15] Tijeni Delleji, Feten Slimeni, Hedi Fekih, Achref Jarray, Wadi Boughanmi, Abdelaziz Kallel, and Zied Chtourou. An upgraded-yolo with object augmentation: Mini-uav detection under low-visibility conditions by improving deep neural networks. *Operations Research Forum*, 3(60):1–27, 2022. 1, 2, 3
- [16] Naqqash Dilshad, JaeYoung Hwang, JaeSeung Song, and NakMyoung Sung. Applications and challenges in video surveillance via drone: A brief survey. In *2020 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 728–732. IEEE, 2020. 1
- [17] Di Feng, Ali Harakeh, Steven L Waslander, and Klaus Dietmayer. A review and comparative study on probabilistic object detection in autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):9961–9980, 2021. 3, 6
- [18] Dario Floreano and Robert J Wood. Science, technology and the future of small autonomous drones. *nature*, 521(7553):460–466, 2015. 1
- [19] Gianni Franchi, Xuanlong Yu, Andrei Bursuc, Angel Tena, Rémi Kazmierczak, Séverine Dubuisson, Emanuel Aldea, and David Filliat. Muad: Multiple uncertainties for autonomous driving, a benchmark for multiple uncertainty types and tasks. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. 1
- [20] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. 3
- [21] Dayan Guan, Jiaying Huang, Aoran Xiao, Shijian Lu, and Yanpeng Cao. Uncertainty-aware unsupervised domain adaptation in object detection. *IEEE Transactions on Multimedia*, 24:2502–2514, 2021. 1, 3
- [22] Ali Harakeh, Michael Smart, and Steven L Waslander. Bayesod: A bayesian approach for uncertainty estimation in deep object detectors. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 87–93. IEEE, 2020. 3
- [23] Ali Harakeh and Steven L Waslander. Estimating and evaluating regression predictive uncertainty in deep object detectors. *arXiv preprint arXiv:2101.05036*, 2021. 3
- [24] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell.

- Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. Pmlr, 2018. 3
- [25] Yo-Ping Huang, Lucky Sithole, and Tsu-Tian Lee. Structure from motion technique for scene detection using autonomous drone navigation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(12):2559–2570, 2017. 1
- [26] Nan Jiang, Kuiran Wang, Xiaoke Peng, Xuehui Yu, Qiang Wang, Junliang Xing, Guorong Li, Jian Zhao, Guodong Guo, and Zhenjun Han. Anti-UAV: A large multi-modal benchmark for UAV tracking, 2021. 1, 2, 3
- [27] Rongqi Jiang, Yang Zhou, and Yueping Peng. A review on intrusion drone target detection based on deep learning. In *2021 IEEE 4th Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC)*, volume 4, pages 1032–1039. IEEE, 2021. 2
- [28] Muhammad Asif Khan, Hamid Menouar, Aisha Eldeeb, Adnan Abu-Dayya, and Flora D Salim. On the detection of unauthorized drones-techniques and future perspectives: A review. *IEEE Sensors Journal*, 2022. 1
- [29] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Mcready. A robust learning approach to domain adaptive object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 480–490, 2019. 3
- [30] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Chang-ick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6092–6101, 2019. 3
- [31] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 4
- [32] Fabian Kuppers, Jan Kronenberger, Amirhossein Shantia, and Anselm Haselhoff. Multivariate confidence calibration for object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 326–327, 2020. 3, 5
- [33] Mattia Litrico, Alessio Del Bue, and Pietro Morerio. Guiding pseudo-labels with uncertainty estimation for source-free unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7640–7650, 2023. 1
- [34] Giulio Mattolin, Luca Zanella, Elisa Ricci, and Yiming Wang. Confmix: Unsupervised domain adaptation for object detection via confidence-based mixing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 423–433, 2023. 8
- [35] Sharifah Mastura Syed Mohd Daud, Mohd Yusmialdil Putera Mohd Yusof, Chong Chin Heo, Lay See Khoo, Mansharan Kaur Chainchel Singh, Mohd Shah Mahmood, and Hapizah Nawawi. Applications of drone in disaster management: A scoping review. *Science & Justice*, 62(1):30–42, 2022. 1
- [36] Mohammad Moshref-Javadi and Matthias Winkenbach. Applications and research avenues for drone-based models in logistics: A classification and review. *Expert Systems with Applications*, 177:114854, 2021. 1
- [37] Muhammad Akhtar Munir, Muhammad Haris Khan, M Sarfraz, and Mohsen Ali. Ssal: Synergizing between self-training and adversarial learning for domain adaptive object detection. *Advances in Neural Information Processing Systems*, 34:22770–22782, 2021. 1, 3
- [38] Kevin P Murphy. *Probabilistic machine learning: an introduction*. MIT press, 2022. 1
- [39] Dang-Khoa Nguyen, Wei-Lun Tseng, and Hong-Han Shuai. Domain-adaptive object detection via uncertainty-aware distribution alignment. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2499–2507, 2020. 1, 3
- [40] Kemal Oksuz, Tom Joy, and Puneet K Dokania. Towards building self-aware object detectors via reliable uncertainty quantification and calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9263–9274, 2023. 3
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 4, 5
- [42] Ágoston Restás. Drone applications fighting covid-19 pandemic—towards good practices. *Drones*, 6(1):15, 2022. 1
- [43] Tobias Riedlinger, Matthias Rottmann, Marius Schubert, and Hanno Gottschalk. Gradient-based quantification of epistemic uncertainty for deep object detectors. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3921–3931, 2023. 3
- [44] Tobias Riedlinger, Marius Schubert, Karsten Kahl, and Matthias Rottmann. Uncertainty quantification for object detection: output-and gradient-based approaches. In *Deep Neural Networks and Data for Automated Driving: Robustness, Uncertainty Quantification, and Insights Towards Safety*, pages 251–275. Springer International Publishing Cham, 2022. 3
- [45] Alejandro Rodriguez-Ramos, Javier Rodriguez-Vazquez, Carlos Sampedro, and Pascual Campoy. Adaptive inattentional framework for video object detection with reward-conditional training. *IEEE Access*, 8:124451–124466, 2020. 1, 2
- [46] Subhankar Roy, Martin Trapp, Andrea Pilzer, Juho Kannala, Nicu Sebe, Elisa Ricci, and Arno Solin. Uncertainty-guided source-free domain adaptation. In *European Conference on Computer Vision*, pages 537–555. Springer, 2022. 1
- [47] Artem Rozantsev, Vincent Lepetit, and Pascal Fua. Flying objects detection from a single moving camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4128–4136, 2015. 1, 2, 3
- [48] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. Beyond sharing weights for deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(4):801–814, 2019. 1, 2, 3
- [49] Chen Rui, Guo Youwei, Zheng Huafei, and Jiang Hongyu. A comprehensive approach for uav small object detection with simulation-based transfer learning and adaptive fusion. *arXiv preprint arXiv:2109.01800*, 2021. 1

- [50] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In *European Conference on Computer Vision (ECCV)*, pages 707–724, 2018. [1](#)
- [51] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021. [1](#)
- [52] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8503–8512, 2018. [3](#)
- [53] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics: Results of the 11th International Conference*, pages 621–635. Springer, 2018. [4](#)
- [54] Javad Shahmoradi, Elaheh Talebi, Pedram Roghanchi, and Mostafa Hassanalian. A comprehensive review of applications of drone technology in the mining industry. *Drones*, 4(3):34, 2020. [1](#)
- [55] Han Sun, Jian Yang, Jiaquan Shen, Dong Liang, Liu Ning-Zhong, and Huiyu Zhou. TIB-Net: Drone detection network with tiny iterative backbone. *IEEE Access*, 8:130697–130707, 2020. [1](#), [2](#), [3](#)
- [56] Tao Sun, Mattia Segu, Janis Postels, Yuxuan Wang, Luc Van Gool, Bernt Schiele, Federico Tombari, and Fisher Yu. Shift: a synthetic driving dataset for continuous multi-task domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21371–21382, 2022. [1](#), [3](#)
- [57] Fredrik Svanström, Cristofer Englund, and Fernando Alonso-Fernandez. Real-time drone detection and tracking with visible, thermal and acoustic sensors. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7265–7272. IEEE, 2021. [1](#), [2](#), [3](#)
- [58] Bilal Taha and Abdulhadi Shoufan. Machine learning-based drone detection and classification: State-of-the-art in research. *IEEE access*, 7:138669–138682, 2019. [1](#)
- [59] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2517–2526, 2019. [6](#)
- [60] Viktor Walter, Matouš Vrba, and Martin Saska. On training datasets for machine learning-based visual relative localization of micro-scale UAVs. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10674–10680. IEEE, 2020. [1](#), [2](#), [3](#)
- [61] Jinyu Yang, Mingqi Gao, Zhe Li, Shang Gao, Fangjing Wang, and Feng Zheng. Track anything: Segment anything meets videos, 2023. [4](#)
- [62] Chunhua Zhang and John M Kovacs. The application of small unmanned aerial systems for precision agriculture: a review. *Precision agriculture*, 13:693–712, 2012. [1](#)
- [63] Jin Zhang, Xi Zhang, Yanyan Zhang, Yexin Duan, Yang Li, and Zhisong Pan. Meta-knowledge learning and domain adaptation for unseen background subtraction. *IEEE Transactions on Image Processing*, 30:9058–9068, 2021. [1](#), [4](#)
- [64] Jie Zhao, Jingshu Zhang, Dongdong Li, and Dong Wang. Vision-based anti-uav detection and tracking. *IEEE Transactions on Intelligent Transportation Systems*, 23(12):25323–25334, 2022. [1](#), [2](#), [3](#)
- [65] Shen Zheng, Changjie Lu, Yuxiong Wu, and Gaurav Gupta. Sapnet: Segmentation-aware progressive network for perceptual contrastive deraining. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 52–62, 2022. [8](#)
- [66] Ye Zheng, Zhang Chen, Dailin Lv, Zhixing Li, Zhenzhong Lan, and Shiyu Zhao. Air-to-air visual detection of micro-UAVs: An experimental evaluation of deep learning. *IEEE Robotics and Automation Letters*, 6(2):1020–1027, 2021. [1](#), [2](#), [3](#)